

Revenue Models, Price Differentiation and Network Neutrality Implications in the Internet

Costas Courcoubetis, Kostas Sdrolias
Department of Computer Science
Athens University of Economics and Business
Athens, GR
courcou@aueb.gr, sdrolias@aueb.gr

Richard Weber
Statistical Laboratory
University of Cambridge
Cambridge, UK
rrw1@cam.ac.uk

1. INTRODUCTION

The long-standing network neutrality debate concerns the future of pricing practices in the Internet [2, 5]. Presently, an Internet service provider (ISP) charges both its end-user subscribers and its directly connected content service providers (CSPs) for access to its network and the whole Internet. A central question is whether this practice should continue, or whether each ISP should be permitted to impose charges upon all content providers that deliver content to the ISP's subscribers (not only upon those that attach to the ISP directly). The debate is thus between the existing practice of 'one-sided pricing' (also called, 'neutral' pricing), in which the ISP obtains all its revenue by only charging its end-users, and 'two-sided pricing' (non-neutral) [1, 9].

There is a substantial literature on the subject, [3, 4, 6, 7], focusing on the welfare effects of two-sided pricing compared to neutral pricing. Many different models have been proposed; using non-comparable sets of parameters and assumptions, they have produced differing answers. In this paper we present a model which is both simple and general, based on the model proposed in [8]. It captures the essential aspects of investment incentives by ISPs and CSPs and how these are influenced by the different ways that ISPs might charge. We allow for the fact that CSPs may be of different sizes and that they can obtain revenue from selling content to end-users (as well as from from advertisement, as in [8]). We simplify the model by focusing on a single ISP who interacts with many independent CSPs. As in [8] we analyse what happens at the symmetric Nash equilibrium of the induced leader-follower game. Our results are not concerned with the welfare effects of adopting alternative pricing policies, but on the incentives of the CSPs to accept them.

We are able to make some interesting observations about the incentives provided to large and small CSPs. If all CSPs have as their business model to sell content to end-users at some strategically defined price, then both large and small CSPs will have aligned incentives towards adopting or refusing a two-sided pricing policy if this is proposed by the ISP; it is either beneficial for all or for none. In implementing two-sided pricing, an ISP might use either a 'uniform' or 'differentiated' pricing strategy; in the latter case the ISP charges each CSP a price that depends on the CSP's size, and in the former case not. CSPs may have differing preferences for these strategies. On the one hand, if all CSPs aim

to make revenue from both advertisement and by charging their customers for content, then no benefit can be obtained by the ISP through using differentiated pricing. All CSPs will be charged the same, and under two-sided will incur the same improvement or degradation of their profits, as compared to neutral pricing. On the other hand, if all CSPs have as their business model to make revenue only from advertisements, then small CSPs will favor two-sided differentiated pricing, over two-sided uniform pricing, whereas large CSPs will have the reverse preference. This suggests that differentiated pricing can be used to impose fairness, since the large CSPs, having greater market power, pay for a larger portion of the infrastructure investments and allow for the smaller CSPs to enjoy some extent of 'free-riding'. This facilitates entry of small CSPs while increasing the total welfare (the welfare effects are not shown in this short paper). Hence the differentiated pricing approach might be preferred by the regulator.

In what follows we define the economic model and the key parameters. We conclude with the key results of the analysis, omitting the proofs and calculations which will appear in the full version of the paper.

2. THE ECONOMIC MODEL

We consider a scenario in which a single large ISP offers access services to customers, and interconnection to a number of content service providers (CSPs) who sell services to the access customers of the ISP. The CSPs obtain additional revenue by renting space for ads on their portal (or by incorporating ads in their service interfaces). We reasonably assume that this revenue is of the form $\rho \times a$, where ρ is the rate of visits to a site, i.e. the click rate that is due to the popularity of a site's content, and a is the normalized ad revenue per user click. We may also view a as the charge per click that advertisers pay to a site, calculated as the quotient of the total advertisement budget allocated to that site divided by the number of clicks the site obtains. Different sites may be able to charge different ad prices due to their popularity, type of users visiting the sites, etc.

We consider two types of CSPs: type 1 who charge advertisers a fee of a_1 per click and type 2 who charge $a_2 = a_1/2$ per end-user click. We shall often refer to types 1 and 2 as 'large' and 'small', respectively, since in many practical instances the larger content sites are the ones that tend to charge greater ad rates. We investigate if, and under what market conditions, any type of player in the ecosystem might benefit if the access ISP were allowed to charge the CSPs.

Additionally, we examine under what circumstances the exact market power of the CSP determines the pricing scheme that most benefits it. To obtain more detailed answer to these questions we study two different end-user charging schemes which reflect current practice in today's Internet.

But first we need to explain some aspects of our model which make it simple and consistent. As in [8] we define all the relevant economic measures as functions of the click rates of the users. We suppose that all the access traffic is due mainly to browsing sites for content. Hence the ISP's traffic can be defined in units of click rates. Similarly, the revenue that CSPs obtain from selling content (videos, news, applications, etc), can also be expressed as revenue per click, by forming the quotient of the total revenue and the total number of clicks to the site. Again, one can interpret this as an average price per click charged by the CSP to its customers. This price is zero if the CSP obtains revenue only from advertisement. Two models are precisely defined as:

- Case 1: End-users face a fixed price of $s + p$ per click, where s is charged by the access ISP and p is charged by the CSP (of the page on which the click is made).
- Case 2: End-users face only a fixed price of s per click, where s is charged by the access ISP.

In both cases, the price s is defined by the market competition and conditions and is exogenous to our model and not part of the strategy space of the access ISP. For instance, in current practice many ISPs charge flat rates and $s = 0$ (no variable part in the tariff). If users pay usage charges when consuming above a given quota (say, 40Gb download per month) then $s > 0$ is sensible. We can also handle within our model the scenario in which a high quality service (HQS) is deployed by a CPS and end-users are charged more for it than for a lesser quality service. Similarly, CSPs might wish to develop content which to be successful needs the ISP to provide a high quality access service. Or their current offerings might benefit from higher quality ISP services. Now the access ISP might charge both its own end-user customers and the CSP who is acquiring the HQS special fees.

Our model can also handle the circumstance in which the ISP attempts to capture all of the additional revenue earned by a CSP, via a two-sided pricing scheme, and while also charging his access customers as before. There are many other scenarios that can be represented by our model.

A key aspect of the model (following [8]) is the inclusion of incentives for investments by the ISP and the CSPs. We suppose that the rate at which a typical customer of the ISP purchases the content of any CSP (e.g. films watched) depends on the quality investments of the access and the content provider, together with the charges imposed by both types of providers. Both ISP and CSPs expect greater traffic (and hence revenue) when the qualities of the transport and content services increase. Let t denote the level of the ISP's investment in the HQS, and let c_1 and c_2 denote, respectively, the investment in content quality by each large and small CSP. By converting the price paid for access traffic and for purchasing content on a per click basis, we suppose that the purchase rate of the content of each type i CSP is:

$$\rho_i = c_i^u t^w e^{-(s+p_i)/\theta},$$

where $\theta > 0$ and $w, u \geq 0$ with $w + u < 1$. The parameter θ denotes the end-user price sensitivity, while u and w are

parameters which control the shape of the demand for these services as a function of these investments. Normally, it is sensible to take $w < 1$ and $u < 1$, so that ρ_i is a concave increasing function of both c_i and t . Case 2 is characterized by the restriction that $p_1 = p_2 = 0$.

We now describe the distinction between one-sided and two-sided pricing. In one-sided pricing the ISP obtains all its revenue from the end users. Supposing that the numbers of CSPs of types 1 and 2 are n_1 and n_2 , respectively, the total number of clicks from the end-users is $n_1\rho_1 + n_2\rho_2$, and thus with one-sided pricing the ISP has revenue of $(n_1\rho_1 + n_2\rho_2)s$. (We analyze only the part of the ISP's revenue that is relevant to our model; it may have other revenues from other sources). In two-sided pricing the ISP charges a price of q per click to the CSPs for giving them the right to have access to its customer base. Given that the CSPs are not identical in terms of their market power, it is also interesting to investigate the implications if the access ISP may charge different types of CSP differentially.

Each provider's objective is to maximize its profit, that is, the revenue generated by the end-users clicks, minus investment costs. For each pricing scheme that we study we find the symmetric equilibria for both one- and two-sided pricing, and compare each providers profits in both regimes. This equilibrium is found within the context of a two-stage leader-follower game. In stage 1 the access ISP (Stackleberg leader) chooses its level of investment t , and the prices it charges to the CSPs (q_1 and q_2) so as to maximize its profits. Subsequently in stage 2 each CSP chooses its corresponding level of investments c_i , and end-user price p_i (if allowed), so as to maximize its profit. We allow both q_i and p_i to be negative to model subsidies. Thus

- ISP profit: $\pi_{\text{ISP}} = (s + q_1)n_1\rho_1 + (s + q_2)n_2\rho_2 - kt$,

where k is a given constant. If price discrimination is not allowed then $q_1 = q_2 = q$, which means that CSPs of any type pay the same price per end-user click to the access ISP. In the one-sided market access ISP cannot charge the CSPs, thus $q_1 = q_2 = 0$.

The profit function of a type i CSP is

- CSP i profit: $\pi_{\text{CSP}_i} = (a_i + p_i - q_i)\rho_i - bc_i$,

where b is a given constant, the same for both $i = 1, 2$. In Case 2 we impose $p_1 = p_2 = 0$, as the CSPs do not charge the end-users and their revenues are generated only by their advertisement rates.

3. SHARING REVENUE WITH CONTENT PROVIDERS

In this section we consider how a two-pricing scheme affects the profits of each player in the ecosystem. We also investigate if, and to what extent, the revenues of the providers are affected if the ISP uses a differentiated pricing strategy (in which $q_1 \neq q_2$). We do this by comparing the symmetric Nash equilibria of the one- vs two-sided pricing in terms of the following ratios:

$$r^{\text{ISP}} = \frac{\pi_{\text{ISP}}(\text{one-sided})}{\pi_{\text{ISP}}(\text{two-sided})}$$

$$r^{\text{CSP}_i} = \frac{\pi_{\text{CSP}_i}(\text{one-sided})}{\pi_{\text{CSP}_i}(\text{two-sided})}, \quad i = 1, 2.$$

We find these ratios for a range of the ecosystem variables, addressing cases 1 and 2 in turn.

3.1 Case 1: CSPs charge end-users

The growth of services such as “TV-ready broadband” and “PayTV” demonstrates that end-users are willing to pay for premium content delivery. By computing the symmetric equilibrium we find that when in case 1 the ISP uses two-sided pricing it does not gain by using a differentiated pricing strategy: it is optimal simply to charge both types of CSP the same price q per click. Specifically,

$$q_1 = q_2 = q = (1 - u)\theta - s.$$

Notice that q does not depend on the advertisement-based revenues of the CSPs or the number of the CSPs of each type. It also turns out that the optimal price to be charged by a CSP is $p_i = \theta + q_i - a_i$. Observe that if a_i is large then a type i CSP wishes to subsidize users to click on its content.




Figure 1: Ratio of profits for access ISP (solid) and CSPs (dotted)

In all of Figures 1–6 we have taken $u = 0.5$, $w = 0.3$, $\theta = 20$. In Figure 1, where we have taken $a_1 = 30$ and $n_1 = n_2 = 40$, the solid curve is a plot of r^{ISP} against s . Notice that at each s we have $r^{\text{ISP}} \leq 1$ so the ISP is always at least as profitable in the two-sided market as in the one-sided market. The ISP is only indifferent if market conditions lead it to set $q = 0$ (despite being permitted to charge CSPs for their transaction with its customer base). In the figure this occurs for $s = (1 - u)\theta = 10$.

Rather surprisingly, the introduction of two-sided pricing increases the profits of both small and large CSPs by the same factor. In the full paper we prove this fact: namely, that $r^{\text{CSP}_1} = r^{\text{CSP}_2}$; moreover, these values are independent of the numbers n_1 and n_2 . It is interesting to know that incentives for two-sided pricing are aligned for small and large CSPs. This is not true if the CSPs obtain all their revenue from advertisement, as we see in §3.2. The dotted curve in Figure 1 plots r^{CSP_i} against s . Notice that for the extreme values of s/θ a CSP is more profitable under a two-sided pricing regime than under a one-sided regime. When s is large it is advantageous that the ISP subsidizes the CSPs ($q < 0$), and when s is small that the CSPs subsidize the ISP. For intermediate values of s CSPs have greater net-benefit in the one-sided market.

3.2 Case 2: CSPs do not charge end-users

In this case the profits of the CSPs are mainly advertisement-based and depend upon the popularity of the specific content. We investigate how a price q (either uniform or differentiated), set by the access ISP, affects the profits of the providers of each type compared to one-sided pricing.

We conclude that the equilibrium price paid by the CSPs to the ISP per end-user click depends on the advertising revenues of each CSP and whether or not the access ISP follows a differentiated pricing strategy. Thus the ratios of

profits for CSPs with different market power are unequal, which contrasts with to what happens when CSPs charge end-users. In our comparisons we consider the profits of each provider as a function of the a_2/θ ($= a_1/2\theta$). We first study the effect on the profits of the providers when the access ISP charges the CSPs differentially, and then when the ISP charges a uniform price q . In the following two subsection (§3.2.1–3.2.2) we suppose that the revenues of CSPs are only ad-based, taking no extra fees from the end-users.

3.2.1 The ISP charges CSPs differentially

Suppose that the access ISP establishes a pricing strategy in which it charges each type of CSP differentially, as a function of its particular market power (advertising rate). To investigate whether or not CSPs benefit if the access provider uses a differentiated pricing strategy, we again study, for a range of ISP access prices and demand elasticity ratios s/θ , the impact of this charging policy on the ratio between providers’ profits in one- and two-sided markets.

In Figures 2–6 we plot the ratio of profits against the value of a_2 and compare what happens for different mixes of numbers of each type of CSP, namely, (a) $n_1 = 60$, $n_2 = 20$ and (b) $n_1 = 20$, $n_2 = 60$. In all of Figures 2–6 we have taken $u = 0.5$, $w = 0.3$, $\theta = 20$. The dotted, dashed and solid lines are for $s = \theta/4$, $s = \theta$ and $s = 2\theta$, respectively.

The access ISP is strictly more profitable under two-sided pricing than the equivalent one-sided, as we always have that $q_1 + q_2 \neq 0$, and hence the plots for the ISP’s ratio of profits are omitted.




Figure 2: Large CSP ratio of profits r^{CSP_1}




Figure 3: Small CSP ratio of profits r^{CSP_2}

Figure 2 (dotted) shows that it is for large values of a_1/θ and small ISP end user charges s that large CSPs are more profitable in the two-sided market. This is expected since by paying part of their profits to the ISP they subsidize it to invest a greater t in network infrastructure. They are more profitable with one-sided pricing when a_1/θ is intermediate, and then again prefer two-sided pricing when a_1/θ is small

since q turns out to be negative (providing a subsidy from ISP which encourages them to invest). For medium (dashed) to large (solid) values of s the large CSPs prefer one-sided pricing for a large range of their parameters.

In contrast to large CSPs, small CSPs benefit more when the access provider uses a differentiated pricing policy (i.e. their r^{CSP_2} is consistently less than or near 1). Figure 3(a) shows that small CSPs are always more profitable under two-sided pricing when $n_1 \gg n_2$, whereas large CSPs may not be (Figure 2(a)). When $n_1 \ll n_2$ we see in Figure 3(b) that it is only for a small range of values of a_2/θ that small CSPs prefer one-sided pricing, and if this is the case, by a very small margin. Also small CSPs are increasingly more favored in the two-sided market as n_1 increases relative to n_2 since the system gets more subsidized by the large CSPs.

For all CSPs, large or small, we see that as s/θ increases the range of a/θ values for which one-sided pricing is more profitable increases; the ISP has sufficient revenue on its own to invest in infrastructure and the main impact of transferring income to the ISP would be to reduce the CSP profits.

3.2.2 The ISP charges all CSPs the same

Now suppose that the access provider is not allowed to charge CSPs differently as a function of their advertising-rates. Instead, it must set a single price q for any CSP.




Figure 4: Access ISP ratio of profits r^{ISP}

Figure 4 shows that in the two-sided market the ISP is always at least as profitable as in the equivalent one-sided, for any n_1 and n_2 , and for any end-user access price s . The only case in which the ISP is indifferent of the pricing regime is when the market conditions force it to set $q = 0$, although it is allowed to charge the CSPs. One can also see that there is a slight difference in the ratio of the access provider's profits for different mixes of numbers of large and small CSPs, but this does not affect the fact that the access ISP benefits from two-sided pricing (but less so as the end-user charge s grows as the ISP then has sufficient revenue to invest in infrastructure and the ability to charge CSPs does not have such a marked effect on its revenues).

The deleterious effect of a uniform price upon the small type of CSP can be seen by comparing Figures 3 and 6. For most values of s and a_2 (and especially when $n_1 \gg n_2$), the small CSPs prefer one-sided pricing, whereas with a differentiated price they prefer two-sided pricing. This is because the uniform price q is computed by targeting the capabilities of the large CSPs and hence it becomes excessive for the small CSPs. These results suggest that small CSPs will, for most ranges of parameters, be opposed to non-neutral regimes (unless they are charged differentially). Conversely, we see from Figures 2 and 5 that the large CSPs are less

opposed to two-sided pricing if it is uniform, since they will then be charged less than if it is differentiated.




Figure 5: Large CSP ratio of profits r^{CSP_1}




Figure 6: Small CSP ratio of profits r^{CSP_2}

This concludes our analysis, and substantiates our claim in the introduction that differentiated pricing can provide an element of fairness for small CSPs and sustainability of the Internet ecosystem.

4. REFERENCES

- [1] M. Armstrong. Competition in two sided markets. *RAND Journal of Economics*, 37(3):668–691, 2006.
- [2] T. Berners-Lee. Net neutrality: This is serious. *timbl's blog*, 43(4), June 2006.
- [3] J. P. Choi and B.-C. Kim. Net neutrality and investment incentives. *RAND Journal of Economics*, 41(3):446–471, March 2010.
- [4] N. Economides and J. Tag. Network neutrality on the internet: A two-sided market analysis. *Information Economics and Policy*, 24:91–104, 2012.
- [5] D. Farber and M. Katz. Hold off on net neutrality. *Washington Post.*, January 2007.
- [6] R. S. Lee and T. Wu. Subsidizing creativity through network design: Zero pricing and net neutrality. *Journal of Economic Perspectives*, 23(3):61–76, July 2009.
- [7] R. T. B. Ma and V. Misra. The public option: a non-regulatory alternative to network neutrality. *ACM Conference on emerging Networking Experiments and Technologies (CoNEXT)*, December 2011.
- [8] J. Musacchio, G. Schwartz, and J. Walrand. A two sided market analysis of provider investment incentives with an application to the net-neutrality issue. *Review of Network Economics*, 8(1):22–39, March 2009.
- [9] J.-C. Rochet and J. Tirole. Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4):990–1029, June 2003.